Science & Global Security, 21:70–92, 2013
ISSN: 0892-9882 print / 1547-7800 online
DOI: 10.1080/08929882.2013.755028

Applications and Limitations of Nuclear Archaeology in Uranium Enrichment Plants

Matthew Sharp

Bureau of European and Eurasian Affairs, U.S. Department of State,
Washington, DC, USA

The uranium-235 content of a uranium enrichment plant's product is related to the uranium-234 content of its waste, allowing one to check with tails measurements consistency with a plant's declared past production. Verification works best with known feed material, but with unknown feed isotopics the production of low and high enriched uranium may still be distinguished based on tails measurements. Estimating product masses is harder, and concealment scenarios are discussed. With traditional nuclear accounting, relationships between product and waste isotopics, or "nuclear archaeology," can increase confidence in the accuracy of declarations of past fissile material production.

Nuclear archaeology seeks to develop a suite of methods that can provide independent verification or cross-checks of the past production of plutonium and highly enriched uranium. Such tools could enhance confidence in state declarations of past fissile material production, thereby strengthening confidence in steps towards nuclear disarmament.[1] Archaeology in the context of plutonium production requires the measurement and modeling of activation products in the structure of a nuclear reactor in order to estimate the total integrated neutron fluence and, hence, the total amount of plutonium produced. This

Received 18 June 2009; accepted 6 November 2012.

The author gratefully acknowledges the support of a research fellowship from the Project on Managing the Atom and the International Security Program at Harvard University's Belfer Center for Science and International Affairs; the work presented here was carried out under this generous support. The views expressed in this article are those of the author, and do not necessarily reflect those of the U.S. Department of State or the U.S. Government. Helpful discussions and comments are also acknowledged from Frank von Hippel, Steve Fetter, Alex Glaser, Matthew Bunn, David Wright, Hui Zhang, Martin Malin, and Elena Rodriguez-Vietez.
Address correspondence to Matthew Sharp, PhD, Bureau of European and Eurasian Affairs, U.S. Department of State, 2201 C st. NW, Washington, DC, USA. E-mail: matthew.k.sharp@gmail.com

technique has been carefully studied and has been successfully applied in a test case.[2]

This article explores one way in which nuclear archaeology could be applied to uranium enrichment operations. There is a relationship between the enrichment of uranium-235 in the product of an enrichment plant and the depletion of the isotope uranium-234 (present at low levels in natural uranium) in its waste, or its tails. This relationship implies that measurements of the tails can be useful in confirming a declared production history.[3]

This analysis assumes that a country has produced a complete reckoning of its production of enriched uranium, and the aim is to confirm that the country has not concealed the production of a militarily significant quantity of highly enriched material from this declaration. The analysis proceeds under the assumptions that the inspector has access to the declaration, including a chronology of the amounts and enrichment level of all material produced in the plant, and that the waste tailings of the plant's operation are available for inspection, as well as at least samples of the declared enriched product of the plant. With this information, the task is to confirm that significant amounts of highly enriched uranium (HEU) material have not been omitted from the declaration.

The following briefly describes the techniques of nuclear archaeology and gauges their effectiveness in the context of three illustrative cases. For the ideal scenario of a small enrichment enterprise (e.g., only hundreds of kilograms of product total and a single enrichment plant) with feed material of known properties, inspection of the tails can provide good determinations of what kind and how much product material was made. Such a case is the best possible scenario, but it is not completely unrealistic; the IAEA in the early 1990s sought to verify South Africa's declaration of a similar production history. In the case of a similarly simple enrichment program with unknown feed properties (more likely for a small program), a distinction between low and high enrichment levels can still be made.

In larger programs, along the lines of the American effort with hundreds of tons of produced material from several facilities, these kinds of determinations are extremely difficult due to the greater complexity of the plant operation and the widespread destruction of the tails themselves by further enrichment. These schematic models are explored to identify where potential measurements can be useful in testing consistency with a declaration. Scenarios by which undeclared production could be concealed can be identified, and one can seek to identify corresponding cross-checks that would require the concealer to alter material or records in consistent ways, making such scenarios more difficult.

The simplicity of the cases considered here is intended to explore the potential for cross checks and to suggest measurements that could be used to build confidence in a declaration. Employing such tests in more realistic scenarios would require significantly more complex models of facilities, and the accuracy possible would be highly case-specific. Further work will be required to determine if the effectiveness of the tests suggested here in real cases would be commensurate with their costs.

Below, the relevant variables are defined for a model uranium enrichment plant and the constraints that lead to connections among those variables are introduced. These connections are then used to demonstrate the utility of nuclear archaeology in the context of an idealized, simple enrichment plant. The complications of a dramatically more complex production effort are introduced so that the continued efficacy of these methods in such a case can be assessed. The possibility of deriving a date of material processing based on tails measurements is discussed, and, finally, the implications of a number of possible means of concealing material production from these techniques are addressed.

URANIUM ENRICHMENT

Enrichment can be considered as a black box process that turns a flow of feed material into a flow of product, which contains a higher fraction of uranium-235 than the feed, and a flow of depleted tails, which contain less uranium-235. In the case of an ideal cascade, no separative work is lost due to mixing of material of different enrichments. In small-scale plants, a limited number of operating stages or other cost and engineering-related factors may dictate that a cascade is non-ideal; with the details of such a cascade an inspector could seek to construct an accurate model, but the resulting analysis would be highly case-specific. The plant has a characteristic separation capacity, related to the rate at which it can apply work to the feed and measured in separative work units per year, where a separative work unit (SWU) has units of kilograms uranium. The uranium will in general have only three isotopic components: the dominant uranium-238, the fissionable and desired uranium-235, and uranium-234, present naturally in small quantities. The enrichment process can then be described with nine variables:

- The mass of material in each of the plant's streams [kg/yr]: M_f, M_t, M_p
- The enrichment of uranium-235 in the streams: x_f, x_t, x_p
- The enrichment of uranium-234 : y_f, y_t, y_p

Here, subscripts f, t, p indicate that a quantity is measured at the feed, tail, or product stream. The plant's separative capacity, C, is sometimes treated as a tenth variable, but it is not independent from the flow masses. In some cases below, the capacity is treated as an additional variable for convenience, but this variable is tied directly to the absolute mass scale of the plant's operation, not an actual independent and additional variable in the process. In what follows, 10 percent is used as a reasonable uncertainty in a plant's capacity. In practice, achieving even this level of certainty could be extremely unlikely, given that the details of a centrifuge or a gaseous diffusion element may well be considered sensitive information, and all details of a plant's operational history are likely to be known only to the state making a declaration. However, the purpose of this analysis is to explore linkages with verification implications among quantities that can be declared and sometimes measured; deriving an enrichment plant's separative capacity from unfalsifiable data to within 10 percent is likely impossible, but confirming that a 10 percent estimate of a plant's capacity is consistent with all intact records at that level may well be. In what follows, this important caveat on estimating plant capacity will be implicit.

The Y-Plant operated by South Africa at its Valindaba site from 1979 to 1990 provides an excellent illustration of this kind of simple facility. This small plant is particularly relevant because it was used over its lifetime to produce both low-enriched uranium (LEU) for reactor fuel and HEU for nuclear weapons. When South Africa joined the Nuclear Nonproliferation Treaty in 1991, IAEA inspectors invested nearly 2 years in an attempt to reconstruct the output of the plant and to verify that it was not possible that enough undeclared HEU had been produced to construct an additional weapon; this is exactly the verification task nuclear archaeology is intended to fulfill. The Y-Plant used natural uranium as feed and had an average capacity estimated at about 15,000 SWU/year. The South African case is invoked solely as an example of a small program where verification proved both challenging and of particular international interest; the details of this real-life example, including the unique enrichment technology used, are not addressed in what follows.

The nine unknowns listed above are not independent of one another. Garza et al. derive a value function for three components of separation:

$$V(x, y) = c_0 + c_1x + c_2y + c_3H(x, y) + U(x, y), \tag{1}$$

where the c_i are all independent constants.[4] H and U are given by:

$$H(x, y) = yR^{-(2k-1)} \tag{2}$$

$$U(x, y) = \left[2x + \frac{2k(y-1)}{2k-1}\right] ln(R) \tag{3}$$

with

$$R = \frac{x}{1-x-y} \tag{4}$$

and

$$k = \frac{\alpha_y - 1}{\alpha_x - 1} \tag{5}$$

The separation gains for different isotopes, α_i, depend on the technology of separation.[5] For a gaseous diffusion plant, $\alpha = \sqrt{352/349}$ for uranium-235 being separated from uranium-238 in UF_6 gas, and so k for separating uranium-234 and uranium-235 is 1.34. Isotope separation in a gas centrifuge can be characterized by very similar values for α and k. For concreteness, the calculations below use the gaseous diffusion numbers, but the techniques could just as easily be applied to centrifuge enrichment plants. Moreover, it must be noted that more sophisticated methodologies can be used to model cascades, particularly in the case of centrifuge enrichment plants or other cases of large separation factors.[6] However, for the purposes of this analysis, the simpler formalism suffices to explore the interconnections among different declared and measurable quantities related to plant operation.

Sticking to the simpler Garza formulation, $V(x, y)$ is the value function of the process; integrating it over the three flows into and out of the plant (weighted by the mass at each stage) will yield the separative work capacity, C. Since the constants are all independent, each of the five terms on the right side can be integrated separately, with the first four yielding zero and the fifth giving C. The first three of these equations are the familiar mass conservation equations of the system:

$$M_f = M_p + M_t \tag{6}$$

$$x_f M_f = x_p M_p + x_t M_t \tag{7}$$

$$y_f M_f = y_p M_p + y_t M_t. \tag{8}$$

Integrating the last two terms yields:

$$M_f H(x_f, y_f) = M_t H(x_t, y_t) + M_p H(x_p, y_p) \tag{9}$$

$$C = M_t U(x_t, y_t) + M_p U(x_p, y_p) - M_f U(x_f, y_f) \tag{10}$$

With nine variables and only five equations with which to constrain them, four more pieces of information are needed to solve the system (thereby revealing the enrichment of the product). At this level of sophistication, there are no more constraints available to relate the nine quantities. Instead, nuclear archaeology depends on determining four of the variables listed above by some other means.

NUCLEAR ARCHAEOLOGY—EXAMPLES OF APPLICATION

In order to explore the details of implementing archaeology in various scenarios, a series of MATLAB routines have been used to encode the constraints listed in the last section. These programs take as input the feed assay (the concentration of uranium-234 and uranium-235 in the feed), the tails

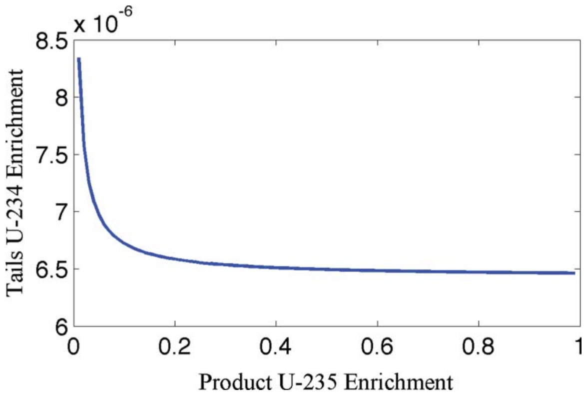

Figure 1: The relationship between the uranium-234 content of the tails and the uranium-235 content of the product for a plant of natural feedstock, $x_f = 0.0072$, $y_f = 5.5 \times 10^{-5}$, and $x_t = 0.002$ (color figure available online).

uranium-235 content, the plant's capacity, and the product uranium-235 enrichment; the routines then give as output the tails assay and all masses. These programs can then be used to generate a curve like that in Figure 1, relating the product enrichment to the uranium-234 content of the tails. This curve is uniquely defined by a plant's feed assay and the uranium-235 content of the tails.

In its simplest implementation, the proposed method of reconstructing uranium enrichment history from the associated waste products depends on the inspector knowing x_f, y_f, and x_t in order to generate such a curve; measuring the uranium-234 content of the tails then allows one to deduce the corresponding product enrichment. Acquiring these three pieces of data should be possible in many cases. One could then measure the tails uranium-235 enrichment, x_t; if the feed is natural uranium, then the uranium-235 content, x_f, is known to be 0.71 percent. The uranium-234 content of natural uranium can in fact vary over some range, as is discussed below, but for the moment it can be considered to be fixed at its average value of 5.5×10^{-5}. This section introduces progressively greater levels of realism and complexity to this scenario and explores how effectively nuclear archaeology can uncover the enrichment and mass of product material as the assumptions are relaxed. For the moment efforts to conceal production are ignored, an issue addressed in a later section.

A Simple Plant with Known Feed Assay and Measurement Errors

A simple uranium enrichment program can be defined as one in which the enrichment plant has a single feed input and two outputs: the tails and the product. If the concentrations of uranium-234 and uranium-235 in the plant's feed material are known, then measuring the assay of the tails and following

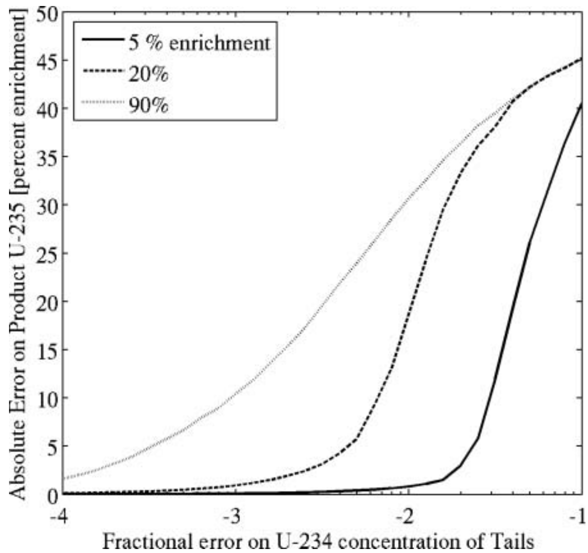

Figure 2: The error on our determination of the uranium-235 content of the product (in percent enrichment) as a function of the fractional measurement error of the uranium-234 in the tails for three levels of product enrichment. Tails uranium-235 content is assumed to be 0.2 percent.

the prescription above determines the assay of the product. In this section, it is assumed that all of the tails are available.

Assuming that the feed material properties are known exactly, the only source of uncertainty in the determination of the assay of product corresponding to a given sample of tails is the random error associated with the measurement of the tails assay. In the following simulations, the relationship between measurement errors in the tails assay and uncertainties in the determination of the product assay for various enrichment levels are explored; the results are shown in Figure 2. Here it is clear that a 1 percent measurement error in the tails implies a$\sim$1 percent error on the determination of product enrichment for LEU product, but that better measurements are required to constrain the enrichment of HEU product to within 10 percent. At the relevant concentrations, techniques like Thermal Ionization Mass Spectrometry (TIMS) can easily produce fractional errors in the range $10^{-4} - 10^{-3}$. Since the errors are random, repeated measurements will increase the accuracy of the measurement. Thus, accurate calculation of the product assay is quite achievable if the feed isotopics are known.

Discovery of HEU Production

There are two scenarios in which nuclear archaeology through measurements of uranium-234 in plant tails can provide verification information. The first is in the case that a country has declared that it has produced only LEU product. The verification challenge in this case is to confirm that no HEU was in fact produced. As shown above, if the inspector has access to the details of the feed and tails of the operation, calculating the product enrichment to within a few percent is possible, and the detection of HEU production would be straightforward. If no HEU production was declared for a given time period, but the tails tell another story, then the attempt to conceal undeclared production will be discovered.

Determination of HEU Product Mass

If the production of HEU was declared, then a second task for nuclear archaeology would be to reveal an underreporting of how much of this product was produced. Having assumed, measured, or calculated the assay of the feed, tails, and product, one can solve Eqs. 6 and 7 for the ratio $\frac{M_p}{M_t}$. With 0.1 percent assay measurements of the tails, one can determine this ratio to within 13 percent for material enriched to 90 percent uranium-235, or 1 percent for material enriched to 4 percent uranium-235. With a measurement of M_t, the mass of product associated with these tails can be determined with the same uncertainties. If the capacity of the plant is known to 10 percent, then the total production of this material over a given time span (even material for which tails are unavailable), can be estimated to the same accuracy.

Determination of Feed Assay

While the above has assumed knowledge of the uranium-234 content of a plant's feed, this isotope's concentration in natural uranium ore can vary from mine to mine and may well be unknown. If the plant drew all of its feed from a single source, then a sample from this source, either as raw ore or as unused UF_6, could reveal its uranium-234 content. However, in 1980, the year that the comparatively simple example of the South African program began operation, more than 20 domestic mines produced 6,000 tons of uranium, from which 200 tons of UF_6 were fed into the enrichment plant.[7] The exact mixture of ore being fed into the plant is likely unknown and inconstant over the plant's operation. In the absence of good data about the feed, an inspector could turn instead to the product. In the South African case, most of the declared product was made available to inspectors. The product, if it can be matched to corresponding tails by dating both, can be used to determine the feed assay. Determining the age of the material depends on measuring the buildup of decay products, which may be removed in any chemical processing that follows enrichment. While the tails will likely not have been chemically treated since leaving the plant, this may not be the case for the product, which is often oxidized or turned into metal before use. This processing may make it impossible to determine the date the material left the cascade, although it would still be possible to determine the date of the material's most recent chemical processing. In some situations, this may be sufficient to determine the feed assay as a function of time.

A Simple Plant with Unknown Feed Assay

In some cases, none of these means of determining the feed assay will be possible, and so one must consider the application of these verification tests in a case in which the uranium-234 content of the feed is unknown. uranium-234 is the result of a decay chain beginning with uranium-238; some of the intermediate stages of this chain can be removed from ore in the presence of moisture, changing the relative abundance of the two isotopes in a way that depends on the local conditions. This results in uranium-234 enrichment levels that can vary across some range from mine to mine. The International Union of Pure and Applied Chemistry (IUPAC) gives this range as $5.0 \times 10^{-5} - 5.9 \times 10^{-5}$.[8] However, recent analysis of samples from many mines has suggested that the range could be substantially smaller, and that natural uranium contains a fraction of uranium-234 within the narrower range of $5.1 \times 10^{-5} - 5.4 \times 10^{-5}$.[9] With this wider range, reconstructing the product enrichment with any accuracy will prove impossible. However, it may still be possible to distinguish the production of material enriched to weapons-grade levels from that destined for use in reactor fuel.

Determination of Product Enrichment

Two means of using the tails to uncover the production of material that is highly enriched without knowing the exact composition of the feed may be suggested. The first relies on assuming a value for the uranium-234 enrichment of the feed; if the possible range of actual values is small enough, this turns out to be sufficient to separate the production of material which is highly enriched from more modestly enriched product. A second method hinges on using the mass of the tails and an estimate of the plant's capacity to solve for the product enrichment. Using the two methods together should allow the successful determination of the broad level of enrichment of the product.

Using an Assumed Feed Assay

Without knowing the precise feed concentration of uranium-234 it is not possible to determine the exact enrichment of the product from the tails alone. However, if the range of possible uranium-234 values is narrow enough, one can often make a guess as to the uranium-234 concentration of the feed and then compute an estimate of the product enrichment close enough to reality to determine whether the actual product enrichment fell above or below a given threshold.

In what follows, tails are simulated to correspond to the production of both 4 percent and 93 percent-enriched material in order to test the ability of the proposed nuclear archaeology methods to distinguish the two under these non-ideal circumstances. In both sets of simulations, the actual feed uranium-234 assay, y_f, is drawn uniformly from a specified range. and the enrichment process is modeled to generate simulated tails. From these simulated tails, the product enrichment can be inferred, using an assumed feed uranium-234 content, y'_f. The value of y'_f that used to reconstruct the product from the simulated tails is in general not the same as the value of y_f that was in fact used to generate these tails. The inferred product is labeled as HEU_t if the reconstructed product enrichment exceeds a threshold t, and as LEU_t if the computed enrichment falls below t. By asking the fraction of tails from the production of 93 percent-enriched material that are correctly identified as HEU_t and the fraction of 4 percent enrichment tails that are correctly identified as LEU_t, the effectiveness of the test can be assessed as a tool for distinguishing the two types of enrichment activity based on the tails and without the exact feed assay.

Two types of errors must be considered in applying this methodology: false positives, where tails from the production of 4 percent material are mistakenly judged to correspond to the production of material more enriched than the threshhold, and false negatives, where tails from to 93 percent-enrichment are mistakenly associated with product less enriched than t. Figure 3 shows how the rate of each class of mistake depends on the threshold. For this test, the value of y_f used to generate the tails is selected randomly from either the IUPAC or Richter ranges, but the feed uranium-234 concentration assumed in the reconstruction, y'_f, is fixed at the center of the allowed range. If the

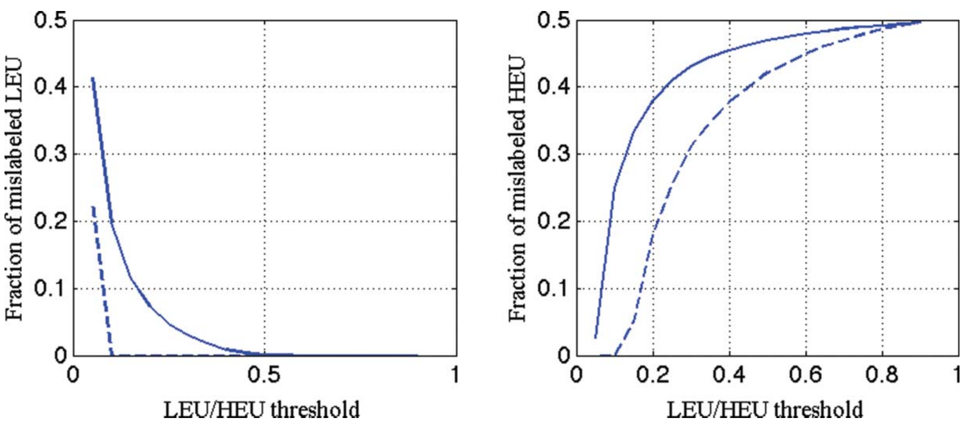

Figure 3: Simulated tails are generated from 4 percent and 93 percent enrichment, picking a value for the uranium-234 content of the feed (y_f) uniformly from either the IUPAC range of allowed values (solid curves) or the narrower range reported in Richter (dashed). Product enrichment is then reconstructed based on these simulated tails, computing a product uranium-235 enrichment by using the center of the allowed range as the assumed uranium-234 concentration in the feed (y'_f). If this computed enrichment exceeds the threshold t, the product is labeled as HEU_t, and if the computed enrichment is less than t then it is labeled as LEU_t. The left panel shows the fraction of tails from 4 percent enrichment that were mislabeled as HEU_t as a function of the threshold t. The right panel shows the fraction of tails from 93 percent enrichment that are mis-reconstructed as LEU_t. These data are the results of inspecting 5,000 sets of simulated tails and include measurement errors (color figure available online).

threshold to separate HEU_t from LEU_t is set to 10 percent enrichment, and the range of possible uranium-234 values is the narrower one specified by Richter, then both the false positive and negative rates go to zero—one can correctly identify whether tails correspond to LEU_{10} or HEU_{10} product 100 percent of the time.

In the case that natural uranium's uranium-234 content varies over the wider range specified by the IUPAC, the same 10 percent threshold mislabels tails from 4 percent material as HEU_{10} for nearly one quarter of possible uranium-234 concentrations, and tails from 93 percent enrichment are mistakenly labeled as LEU_{10} in roughly a quarter of possible scenarios. Note that this is not the probability of incorrectly assessing the product enrichment for a given set of tails, but rather the percentage of all possible enrichment scenarios for which the determination will be wrong. As Figure 3 makes clear, one can adjust the threshold to trade false positives for false negatives in a way that keeps their sum roughly constant. Figure 4 shows which cases are being mislabeled, and how the choice of y'_f can trade false positives for false negatives.

If the range of possible uranium-234 concentration in the feed can be restricted to a range that has the width of the gray band in Figure 4, then a value of y'_f can be assumed such that the production of 4 percent and 93 percent-enriched material can be distinguished using the tails assay. If natural uranium falls within the range given by Richter, then nature has already done this

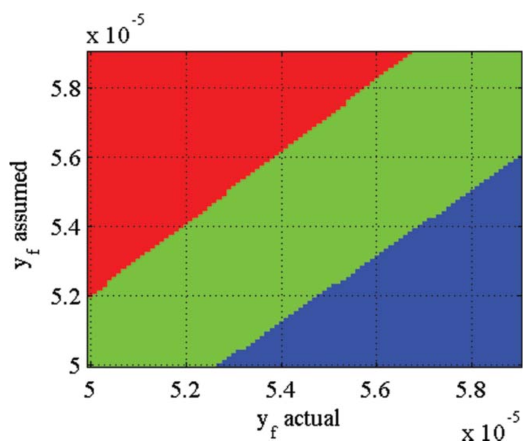

Figure 4: Each pixel represents a combination of the actual y_f (the feed uranium-234 enrichment used to generate simulated tails) and y'_f (the value assumed in reconstructing what kind of product was produced). The central gray region indicates combinations that correctly identify tails from 4 percent enrichment as LEU_{10} as well as tails from 93 percent enrichment as HEU_{10}, the black combinations misidentify the 93 percent product, and the white region misses the 4 percent product. If the possible range of y_f can be limited to about half of the IUPAC range (i.e., closer to the Richter range), there is an assumed value of y'_f that will always generate the correct answer. Within a given range, adjusting the assumed y'_f can trade false positives for false negatives (color figure available online).

work (the allowed range is already narrower than the gray band), and perfect separation is possible for all allowed feed compositions. If the natural range of uranium-234 in uranium ore follows the larger IUPAC numbers worldwide, it may still be true that the range of feed used in any given plant falls within a narrower range. Samples of either ore or product, or even the most general information on where feed was likely to have originated, can help to constrain this range. If the range can be constricted far enough, then Figure 4 will reveal a value for y'_f that a nuclear archaeologist can use to reliably distinguish high and low-enriched tails. If the range remains too broad for this, the plot shows an assumed value that will trade higher false positive rates in return for correctly identifying all 93 percent product as HEU_{10}.

Using the Tails Mass

In many cases, the production of 4 percent and 93 percent-enriched product can thus be distinguished with only knowledge of the range of possible uranium-234 assays in natural uranium feed; in those cases where tails cannot always be correctly associated with LEU or HEU product, one can still arrange to detect production of 93 percent material, incurring higher false positive rates in the process. The tails mass can then be used to separate these false positives from actual weapons-grade material production. This technique relies on the correlation between the mass of tails produced in a fixed time span and the product enrichment, given a fixed capacity and tails assay (panel 1 of Figure 5). With perfect knowledge of the capacity and a good measurement of the tails mass, this correlation yields the product enrichment.

Unfortunately, the enrichment capacity of the plant may not be accurately known. In this event, records of data such as power usage of a plant (available in the South African example) or the number and type of separative units in operation over time could help to form an estimate of the capacity of a plant, and data on feed, tail, and product mass flows at various points in time could be used to calibrate this estimate at those times. If the actual capacity can be estimated to within 10 percent, the distribution of tails masses for 4 percent-enriched and 93 percent-enriched product are shown in panel 2 of Figure 5. The fact that the two regions are non-overlapping shows that, even with an inaccurate estimate of the capacity, the production of 4 percent and 93 percent-enriched material can be distinguished by using the mass of tails created in a given time period.

Measuring the mass in a given tank of tails is easy enough, but determining the production rate of tails requires one also know the time period over which a given tank was filled; this entails another source of uncertainty in any effort to separate the production of HEU from that of LEU. For example, if a 10 percent Gaussian error in the measurement of the tails mass measurement is added to a 10 percent uncertainty in the capacity, there will be significant overlap in the distribution of measured tails masses from the production of

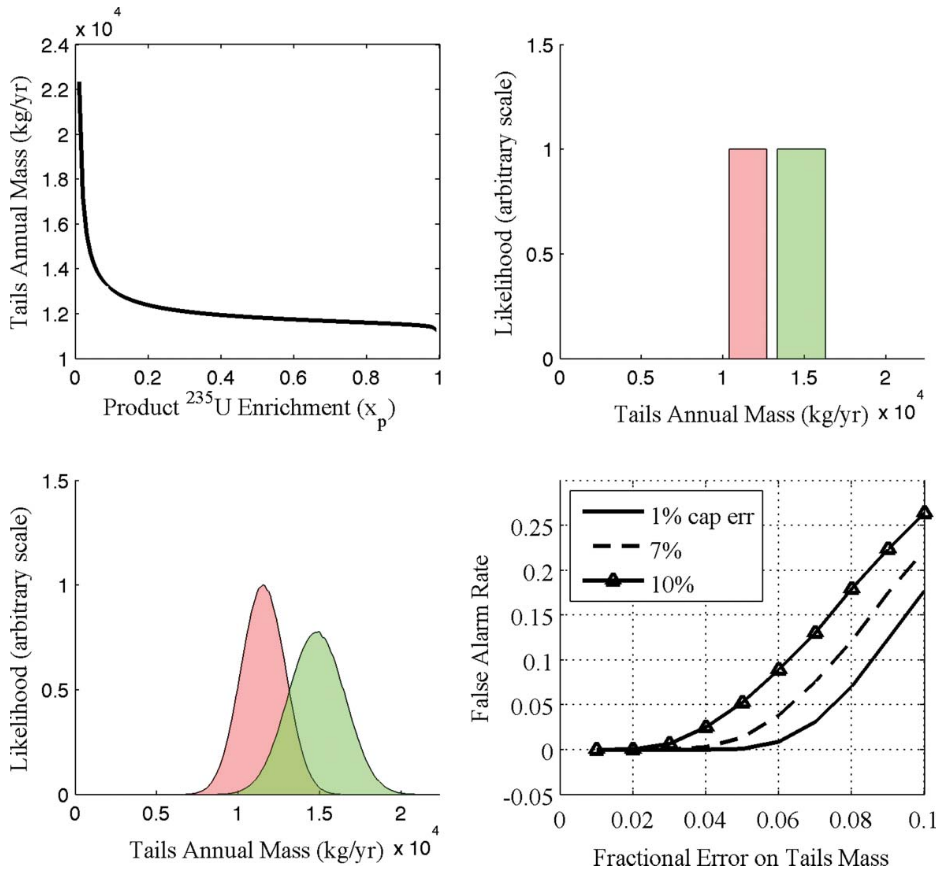

Figure 5: (top left) The relationship between the mass of tails produced each year and the enrichment of product for a 15,000 SWU/year plant operating at 0.2 percent uranium-235 tails. (top right) The distribution of possible mass rates corresponding to 93 percent enrichment (dark) and 4 percent enrichment (lighter), assuming 10 percent uncertainty in the plant's capacity. The two distinct regions indicate that one can still use the mass of tails to distinguish 4 percent and 93 percent enrichment. (bottom left) If a 10 percent error on a measurement of the tails mass is added to the uncertainty in capacity, it becomes more difficult to distinguish the two. (bottom right) The false alarm rate for mis-labeling 4 percent production, given a range of capacity uncertainties and errors on tail mass measurements, keeping the detection probability of 93 percent-enriched material production fixed at 95 percent (color figure avialable online).

reactor and weapons-grade material (panel 3 of Figure 5), making it difficult to reliably distinguish the two by this method. Panel 4 explores how badly this overlap hurts the chances of success; while demanding 95 percent confidence of detecting any enrichment to 93 percent uranium-235 the false positive (mis-identifying tails from 4 percent product) rate is calculated as a function of both the capacity uncertainty and the error on the mass rate. Limiting the false positive rate to the 5 percent level is seen to require ~5 percent mass estimates unless the uncertainty on the capacity can be reduced.

An inspector's ability to meet this goal depends on the details of the plant being inspected and its operation. The contents of a given tank can be weighed and dated accurately; for a 20 year-old program, the mean age of material in a tank can easily be dated to within ~1 week by measuring decay products. If a program produced many containers of tails each year, and if all of these tails are assumed to be available, then an inspector could date all of the tails and sum the masses of those that fall within the desired period. The total mass sent to tails in one year by a plant that was producing one 12-ton container of waste a month or more (about ten times larger than South Africa's Y-Plant) could be easily measured to within 5 percent in this way. If the rate had to be derived from a partial set of tails, or if the production rate for a smaller period of time were desired, then this uncertainty would rise proportionally.

For a smaller plant, fewer waste containers would be filled each year, and a tank with a mean date within a given time window might very well contain substantial amounts of tails from dates outside of this window. In order to determine the tails mass rate in such a scenario, the inspector would need to determine the dates at which the filling of each tank began and ended. Uncertainty on the mass rate would be equal to the uncertainty of these dates, and so 5 percent mass rates would depend on knowing the dates to within a few weeks for a South African-sized plant. Operation records could certainly provide this information, to the extent that these records could be trusted. Alternatively, if the mean ages of two consecutively filled tanks are known, and the filling rate is the same for both tanks, then the transition date is halfway between the two. Where filling rates changed in a known way due to shifts in product assay or plant operation, this can be easily accounted for; shifts in product assay are rare for the small plants considered here. While the time span of tank filling could be misestimated through an unaccounted for change in filling rate or a falsified record, moving the start or end dates will exaggerate the filling rate of one tank by reducing the implied rate of the neighboring tank; an error in one direction is thus balanced by a corresponding error in the opposite direction. In sum, for a variety of plant operation scenarios, 5 percent uncertainties on a measurement of the plant's tails mass rate may be possible, allowing the separation of high and low enrichment. Where such accurate determinations of mass rates are impossible, higher false positive rates would result.

This illustration attempts to distinguish the production of 4 percent and 93 percent-enriched product. Distinguishing the production of two types of material of less different enrichment levels would be more difficult, although the flatness of the tails mass-product enrichment curve above 20 percent enrichment indicates that in most cases it should be possible to distinguish the production of LEU and HEU.

Furthermore, while the tails mass can be altered artificially by hiding a certain quantity of tails, it is only possible to decrease the tails mass in this way. As is clear in Panel 1 of Figure 5, hidden tails could disguise the production of 4 percent-enriched material as that of 93 percent HEU, but not the other way around. If tails from LEU production have in fact been diverted for other uses, such as depleted munitions or recycling, it will be necessary to confirm the quantities involved and to include this material in the accounting in order to avoid the mistaken conclusion of HEU production; diversion of tails for any reason makes the job of a nuclear archaeologist more difficult.

The above presents two means of using the tails to distinguish the production of reactor grade and weapons-grade material. Each of these methods can make this distinction on its own if circumstances are somewhat favorable: if the uranium-234 content of the feed can be somewhat constrained for the first test and if the dating of the tails is sufficiently accurate for the second test. If neither of these criteria is met, then each method will yield unacceptably high false positive rates, but the two methods can still be used in conjunction with one another to discover any tails that correspond to the production of HEU product. A lack of detailed feed isotopics information therefore does not completely eliminate nuclear archaeology's usefulness through the measures discussed here.

Determination of HEU Product Mass

For a collection of tails corresponding to the production of both 4 percent and 93 percent-enriched product, the above tests could be used to isolate all of the tails that correspond to the higher enrichment level. However, without knowing the details of the feed material, the exact enrichment of the HEU product would remain unknown, and it would not be possible to deduce the total product mass from the tails alone. However, one could check for consistency of the tails with the declaration. Because the tails can be dated, they can be matched to specific periods of declared operation. The declared enrichment of these periods determines the ratio of product to tails masses, and the consistency of this ratio with the declared product mass, the measured tails quantity, and the estimated plant capacity will provide increased confidence in the declaration's veracity. While the product mass cannot be determined by these means alone, they do make hiding production that much harder.

LARGER PROGRAMS

Because uncertainty in fissile material production translates into uncertainty in warhead stockpile sizes, an effort to verifiably dismantle the large nuclear arsenals of the United States and Russia would be greatly helped by more accurate assessments of how much HEU was in the possession of both countries. It is unlikely that nuclear archaeology can provide such an assessment through the tools discussed here, due both to the greater complexity of these enrichment efforts and the unavailability of the associated tails.

The simple plants discussed in previous sections involve a single feed and a single product, but the larger programs of the United States and the Soviet Union were more complex than this. In the American case, for example, HEU production was carried out principally at the Portsmouth and Oak Ridge gaseous diffusion plants, but these plants were fed with LEU produced at a separate plant at Paducah, which was in turn often fed with uranium that had been irradiated (and therefore depleted in uranium-235). At the same time, the introduction of uranium-236 into this material adds new constraints on how relative abundances of isotopes change as they move through the system. Sometimes, the plants were operated with multiple product streams, further complicating the accounting.

In principle, detailed records of the processing that led to each sample of tails could be used to construct new relationships between tail and product assays, and one could still determine consistency with a declaration. However, activities in both the United States and the Russian Federation are destroying the evidence contained in the enrichment waste by using these tails for other purposes. In both nations, the tails have been re-enriched in order to extract the uranium-235 remaining in this waste. In particular, more than half of the relevant tails at Paducah were re-fed into the cascades, as were substantial portions of the tails from the other two American plants, often in a mixture of material from different sources. Techniques to make sense of the tails from the reprocessing of such a mix of other tails would involve considerably more ambiguity than those discussed here, and the effectiveness of such techniques would be strongly dependent on detailed records of the various plants' operation. It is doubtful whether such records are available, and they could be easily falsified if they did exist. Nevertheless, such checks of the consistency of declarations would be confidence building and should be investigated further.

The extent to which Russian tails have been re-enriched is less clear, but the total amount of Russian tails still on hand is estimated to be of the same magnitude as the American stockpile. Since Soviet production of HEU exceeded that of the United States, it is likely that an even greater degree of re-enrichment was carried out on these tails. Over the past years, Russian HEU has been blended with 1.5 percent-enriched material to produce LEU to sell to the United States through the HEU deal; while this 1.5 percent material does come from tails-stripping, the tails that are being stripped come from a separate arrangement to re-enrich tails from URENCO, and not from previous Soviet enrichment activity. This particular activity therefore does not impact the application of these measures to the Soviet weapons program.

An additional portion of the stockpile of depleted tails has been used to manufacture depleted uranium munitions. In the mid 1990s, the demand for this material in the United States was approaching 2,000 tons per year, nearly half a percent of the total stockpile of tails. It is possible that up to 10 percent of the total tails available have been removed from the stockpiles for this purpose, and are no longer available for nuclear archaeology, although the removed material does not necessarily correspond to the earlier periods of weapons material production.

Although these activities have made it impossible to fully verify the US and Russian enrichment programs using the tails, the fraction of tails that have not been re-enriched or used for depleted uranium weapons could in principle be used to at least demonstrate consistency with declared activities, provided that they can be accurately dated as discussed above. Unfortunately, even this may not be possible for very long. In 2008 the US completed construction of two plants to convert the remaining half million tons of tails from UF_6 to a more stable oxide form over the coming decades. While such conversion should not adjust the isotopics of the uranium, it would likely involve chemical processing that will make dating the tails very difficult. In programs of such large volumes and ranges of product, an inability to associate tails with a certain production period would be catastrophic for any hopes of using tails for verification purposes; within the next 20 years, any measurements of tails from the American HEU enterprise will be all but impossible. There are no plans to convert the remaining Russian tails from UF_6.

TAILS DATING

The ability to accurately date the available tails is very helpful in determining the mass of tails produced in a given time and in detecting a state that seeks to conceal a diversion of material. Because of the separation processing, the uranium-235 and uranium-238 in the tails is likely to leave the plant clean of any of either isotope's decay products. However, as soon as the tails are removed, each uranium isotope will begin to decay, leaving daughter products at a concentration that depends on the amount of time elapsed. Measuring the concentration of these daughters therefore gives an estimate of when the material was removed from the plant.

For decays with half-lives much longer than the time-scale of the enrichment activity (true of all decay products relevant here), the fractional error in the determination of the age of the material is of the same order as the fractional error in the measurement of the concentration of the decay product. Even at the very small fractional levels involved in uranium decay products, 0.1 percent measurements should be possible in principle with advanced mass spectrometry techniques, if homogenous samples can be acquired.

With 0.1 percent measurements of the age of tails that are less than 10 years old (as in the South African example), the uncertainty is significantly less than the amount of time it takes to fill a tank. In this event, the measured age of a well-mixed tank represents the average age of the tank's contents. Uncovering the actual start and end dates of filling for this tank will require more information, either from records or from the dating of chronologically consecutive tanks, as discussed previously.

While measurement accuracies can be achieved at the levels required for the tests discussed here, real world scenarios will likely involve significant complications. The actual mechanics of acquiring a sample from a container of tails that includes the decay products, whether by physically removing a solid sample of material or somehow acquiring a gas-phase sample that includes both tails material and decay products, may be far from trivial. In this analysis, it has been assumed that a sampling technique is available that allows the measurement of decay product concentrations; the development of such techniques requires further work.

In principle, product material can also be dated in this way, but chemical processing can remove the relevant decay products, allowing only lower limits of age to be put on material. While most product material undergoes some form of chemical processing, the tails, a waste product, are often left alone. As mentioned previously, the US provides an exception to this, having put into action plans to convert all tails to a more stable oxide form. Such treatment dramatically impairs the possibility of dating material.

CONCEALING UNDECLARED PRODUCTION

In much of what has been discussed above, there are obvious ways of altering the record of tails in such a way as to make detection of a diversion of material more difficult for the nuclear archaeologist. Such scenarios can be considered in the case of a plant that mirrors that of South Africa, i.e., one with about 10 years of operation and with the tails, product, and records available for inspection. Consider a scenario in which a weapon's worth of HEU (25 kg) has been diverted and concealed, corresponding to roughly one year of processing time. While hiding or altering tails makes nuclear archaeology through these means impossible, inspectors might use these techniques to complement traditional accountancy to detect inconsistencies in such a scenario.

Erasing Production from Records

If a state attempts to remove the production of 55 kg of HEU from its declaration, but presents the tails from this production to the inspector, then the techniques described here will reveal the production of HEU during the undeclared period and detect the diversion outright.

Hiding Tails and Claiming No Production

Knowing that the tails will reveal the diversion, an aspiring diverter would likely hide them. Because the remaining tails can be dated quite accurately, the gap in available tails will then be detected and an explanation demanded. Perhaps the declarer would claim that nothing was produced by the plant during the year in question.

Evading detection in this scenario would depend on changing a great many records in a consistent way without being caught. In addition to the approximately 30 tons of tails to be hidden away, the declarer must alter the records of the plant's operation to show no activity for an entire year. In the South African example, this would have included doctoring records related to the electricity consumption of the plant, the introduction of ~30 tons of feed material, and the fluoridation of this feed. The state would have to explain the suspicious gap-year convincingly, and hope that interviews with the large number of people involved didn't reveal that anything was awry. Note that the dating of tails allows the inspectors to focus their attention on the relevant year.

Hiding Tails and Claiming LEU Production

Given the scope of falsification needed to claim that no production was carried out in the year for which tails have been hidden, it would perhaps be more sensible to claim that the plant spent the year making LEU product, instead. This scenario avoids many of the problems of the previous one, for example by not requiring an explanation for plant shutdown that doesn't square with any relevant records, and requires changing the quantity of feed (by about 50 percent) instead of hiding it altogether. However, it introduces the question of what happened to the several tons of LEU product that has been claimed but does not exist. Having misplaced both the tails and product for this year would certainly raise suspicion.

Claiming non-Weapons-Grade HEU Production

If the declaration reported the year's production as HEU, but not as weapons ready material (e.g., the 45 percent-enriched material that South Africa used to fuel the Safari research reactor), then the tails would not need to be hidden, so long as the inspector had no hope of finding the exact uranium-234 content of the plant's feed. Although the amount of product material that is missing from the declaration is much smaller (~100 kg) than in the previous scenario, the high enrichment makes this material much more valuable and unlikely to have been lost. Claiming that the material had been burned as fuel would require fairly elaborate misdirection, especially as the entire production of HEU for the Safari reactor was likely only ~70 kg (i.e., small on the scale of the required diversion).[10]

Spreading Diverted Product Across the Plant Lifetime

Instead of altering the records to hide a seventh year of weapons material production, a state could underreport the amount of material produced in each of seven years, hiding 1/7 of the tails from each year as well. This would largely avoid the complication of drawing attention to the year of production that hides the diversion, as in the previous cases. While this scenario is perhaps the most difficult to detect, it does introduce complications of its own. If the product is available and can be dated, then it could be matched to the tails; diverting 1/7 of each year's tails to hide capacity would then require diverting 1/7 of each year's product instead of simply diverting one of seven weapons. Otherwise, mismatched tail and product dates would reveal the diversion. The metalization of HEU might imply that dating the product would yield the age of the metal, and not of the original HEU, but in the case of a material-starved program like South Africa's, in which material was converted to useable forms as fast as it could be produced, the two dates might well be close enough to reveal inconsistencies between the tails and the product.

A further risk to a state seeking to divert material undetected in this way is that the total mass produced each year is fixed by the plant's capacity. If the capacity is known to be 10 percent, the diversion is limited to this level. In South Africa's case, this limits the diversion to ~5 kg of HEU per year for seven years, well short of 55 kg. To hide more material than capacity uncertainties could explain, records of operation would need to be changed to hide the excess capacity, and one might wonder why the capacity was so much lower when HEU was being produced than it was during the rest of the plant's lifetime. Furthermore, assuming South African rates of tank filling, this form of diversion would involve partially emptying tanks of tails. Leaving behind tails tanks that are partially filled might raise suspicion, particularly if the fraction of tails missing and the fraction of capacity unaccounted for matched exactly. This kind of diversion would be easier to hide for a much larger program.

Doctoring Tails

A final option for concealing undeclared production would be to leave the tails in place but to mix extra uranium-234 into them in order to make tails from HEU production look like they came from the production of LEU, instead. By mixing tails from 93 percent enrichment with tails from 1 percent enrichment in a 2:1 ratio, an operator could make it look like the tails came from 4 percent product. However, attempts to date the mixed tails would yield an answer one-third of the way between the actual times of production of the two components. If the tails from 1 percent material were freshly made, this would likely lead to confusing inconsistencies in the chronology of tails. Note also that the tails mass is increased by 50 percent in this example, dramatically altering the implied capacity of the plant or requiring the concealment of this extra material.

If the tails from 1 percent enrichment came from approximately the same time as the HEU being diverted, the diversion would be very difficult to detect. However, the diluting tails would likely need to come from another plant, not part of the declaration under verification to avoid explaining why the 1 percent tails were missing. The concealment becomes fairly convoluted.

In each of these cases, the method of concealing undeclared production is not impossible to get away with. However, these contortions to evade detection do create inconsistencies with other methods of nuclear accountancy, and the effort involved in fooling the inspector is considerable. While concealment is not impossible, it is certainly not easy.

CONCLUSIONS

The above has explored several tests that can be applied to the tails of a uranium enrichment plant to detect the undeclared production of weapons-grade HEU. In cases where natural uranium is used as feed and measurements are possible on the tails:

- Given the uranium-234 content of the feed, the product enrichment can be calculated. If samples of feed or product are available, the uranium-234 content of the feed can be discovered.
- If the product enrichment is known, then the product mass matching a quantity of tails can be computed. To the degree that the capacity is known, the total product mass can be independently computed.
- If the uranium-234 content of the feed is unknown but can be constrained to lie within a narrow enough range, then the production of LEU versus that of HEU can often still be discriminated from the tails. It is possible that the natural variation of uranium-234 concentrations is narrow enough already that this distinction can be made. If, as is also possible, the natural range is slightly broader than what is needed for this to be feasible, it can be narrowed with measurements of feed or product material.
- The mass of the tails, combined with knowledge of the plant's capacity and good enough estimates of the dates during which tails were produced, can offer an independent assessment of whether a plant was producing LEU or HEU. Combining this test with the previous one should allow discrimination in most reasonable cases.
- Without knowing the product enrichment exactly, determining the mass of product produced is difficult, but where a declaration has been made, consistency with the declared mass and enrichment of product can be checked using the tails.
- Concealing undeclared production is possible, but requires elaborate measures to avoid detection.

Provided that an enrichment project is small and simple in design, detailed conclusions can be drawn about the enriched product that corresponds to any available tails. If a large portion of the plant's tails is available, nuclear archaeology can make the concealment of HEU production quite difficult. In this sense, it can provide a valuable verification tool to inspectors of fissile material production complexes. By contrast, the complexity of material flows in a larger enrichment program makes the application of these techniques much more difficult. In this case, even judging consistency between a declaration and available tails could prove challenging, although the possibility deserves further scrutiny.

Clearly, these tools could be helpful in verifying the past production of an enrichment plant after its operation has concluded. In addition, exploring verification techniques can also help illuminate what kinds of data should be taken during a plant's operation to make accounting easier down the road. For instance, records of feed and tails isotopics, including trace isotopes, could lend greater transparency to a plant's operation. Measurements, or even samples, where possible, would of course be preferable to theoretical numbers. In some cases, the tails associated with past production of HEU are being recycled in order to further strip them, making them unusable for many of the tests described here. Where this is true, taking measurements of the tails before they are stripped could prevent this data from being lost forever. Material dating could be facilitated by not mixing wastes from different times, or perhaps even by verifiably sealing containers. Additionally, verification would be made easier if careful records were kept relating to the plant's operating capacity, such as numbers of separative units in operation, power consumption, etc. The capacity of a plant determines the quantity of material produced within it; understanding ways of determining this capacity should be the subject of future work. While such measures might be falsified without detection, they would make concealing material production more complicated and offer greater confidence in a verification project. The earlier in a program that such information can be shared, the more confidence it can inspire.

In essence, many operating parameters of a plant are related to one another in well-defined ways; the more measurements and records that are kept during a plant's operation, the more tests can be performed for consistency with declared activity. By implementing appropriate record-keeping and accounting measures early on, plant managers can facilitate verification later on; program records should be kept with the idea that a program might be audited in the future, in contrast to the way many programs have been operated in the past. In this way, large uncertainties in the past production of fissile materials might be avoided, helping to make verifiable disarmament possible.

At the same time, it must be emphasized again that actual enrichment facilities are immensely more complex than the model facilities considered here.

The conclusions drawn here are based on a number of non-trivial assumptions, including on the possibility of the measurements suggested. Further work is needed to assess how the consistency checks suggested here might be applied, at with what difficulty, in more realistic facilities.

NOTES AND REFERENCES

1. Steve Fetter, "Nuclear Archaeology: Verifying Declarations of Fissile-material Production," *Science and Global Security*, 3 (1993): 3–4, 237–261.

2. Thomas Wood, Bruce Reid, John Smoot, and James Fuller, "Establishing Confident Accounting for Russian Weapons Plutonium," *The Nonproliferation Review*, 9: (2002).

3. Fetter 1993; Drell, S, JASON Program Office, *Verification of Dismantlement of Nuclear Warheads and Controls on Nuclear Materials* (McLean, VA: MITRE Corporation, 1993) 57–58.

4. A. De La Garza, G. A. Garrett, and J. E. Murphy, "Muticomponent Isotope Separation in Cascades," *Chemical Engineering Science*, 15 (1961): 188.

5. H. Wood, "Effects of Separation Processes on Minor Uranium Isotopes in Enrichment Cascades," *Science and Global Security, 16* (2008): 1–2, 26.

6. e.g., E. Von Halle, *"Multicomponent Isotope Separation in Matched Abundance Ratio Cascades Composed of Stages with Large Separation Factors*," in Proceedings of the 1st Workshop on Separation Phenomena in Liquids and Gases, ed. K. G. Roesner, 325–356.

7. M. Ford, "Uranium in South Africa," *J. S. Afr. Inst. Min. Metall.*, Vol. 93, No. 2 (1993), 37–58.

8. K. Rosman and P. Taylor, "Isotopic Compositions of the Elements 1997," *International Union of Pure and Applied Chemistry Technical Report* (1997).

9. S. Richter, A. Alonso, W. De Bolle, R. Wellum, and P. Taylor, "Isotopic 'Fingerprints' for Natural Uranium Ore Samples," *International Journal of Mass Spectrometry*, 193 (1999), 9.

10. David Albright, Frans Berkhout, and William Walker, *Plutonium and Highly Enriched Uranium 1996; World Inventories, Capabilities and Policies*, SIPRI (1997), 387.